Simulation of HED Plasmas (4,050,000 Node hours)

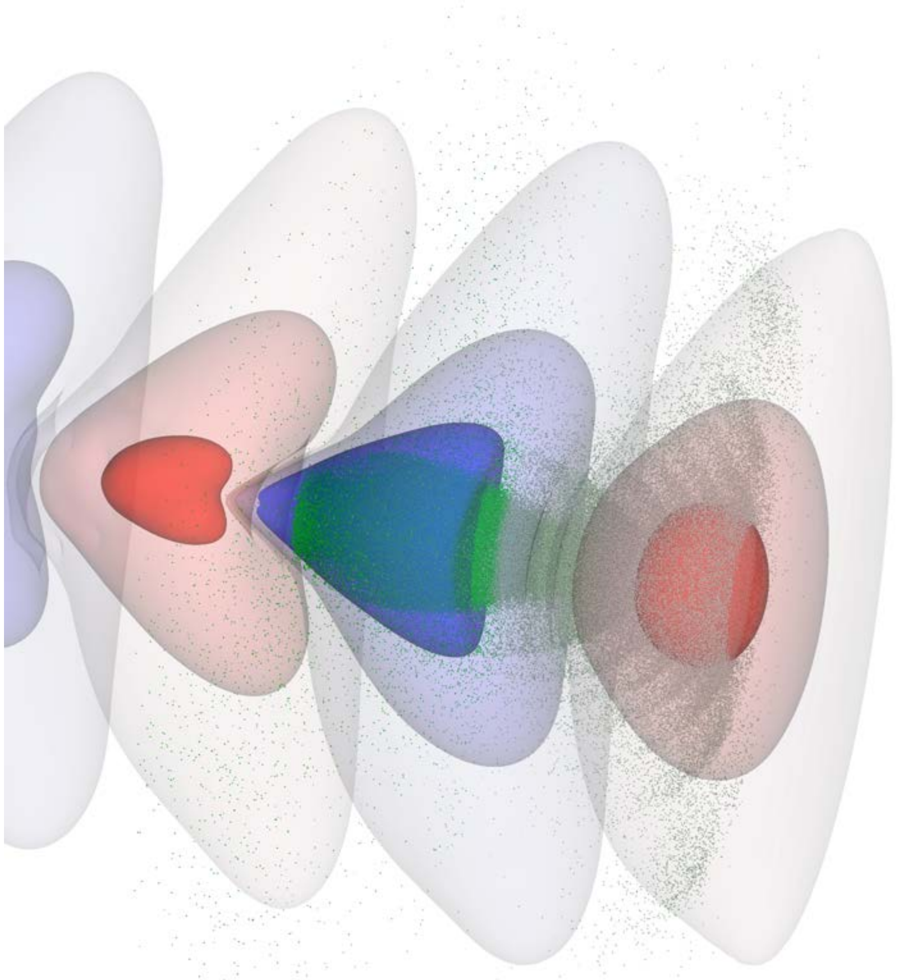

Frank Tsung (co-PI)
Viktor K. Decyk
Weiming An
Xinlu Xu
Han Wen
Thamine Dalichaouch

Warren Mori (PI)
collaborators: L. O. Silva, R. A. Fonseca, IST

PICKSC UCLA

Summary and Outline

OUTLINE/SUMMARY

- Overview of the project
  - **HED plasmas** and the importance of kinetic effects
  - Particle-in-cell method
  - Our main production code — OSIRIS
- Application of OSIRIS to plasma based accelerators:
  - Producing high brightness x-ray using LWFA's.
  - Performing high resolution LWFA simulations in quasi-3D.
  - QuickPIC Simulations of PWFA's.
- Higher (2 & 3) dimension simulations of LPI's relevant to laser fusion
  - Importance of 2D and 3D effects in IFE.
  - Controlling LPI's by temporal incoherence under IFE relevant conditions .
- Code development — porting our codes to the Intel Phi (@ Cori supercomputer @ NERSC), and using deep learning for HED physics.
- Summary/Conclusions

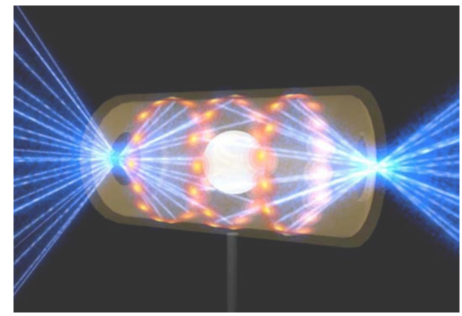

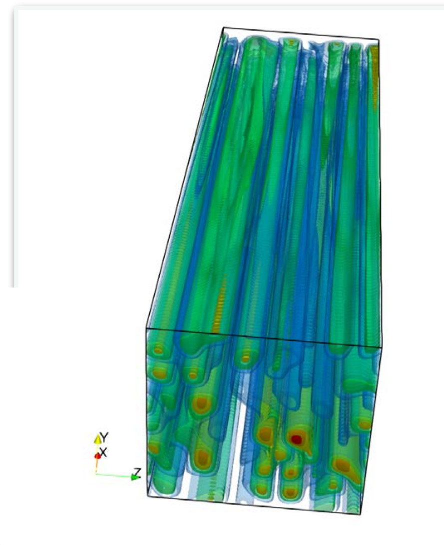

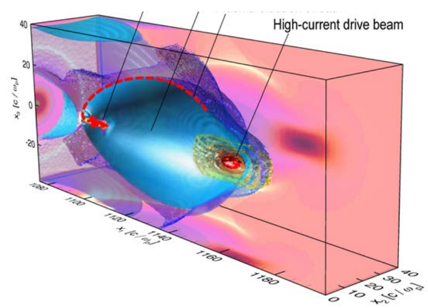

UCLA

osiris framework

- Massivelly Parallel, Fully Relativistic Particle-in-Cell (PIC) Code
- Visualization and Data Analysis Infrastructure
- Developed by the osiris.consortium
  ⇒ UCLA + IST

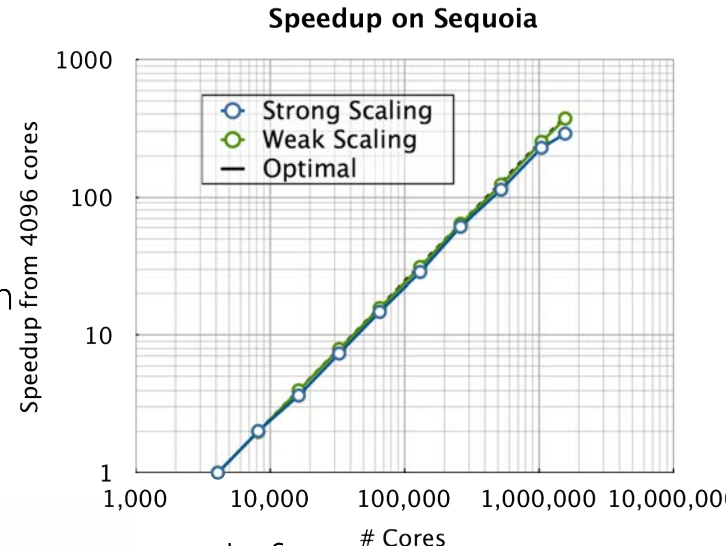

code features

- Scalability to ~ 1.6 M cores (on sequoia).
- SIMD hardware optimized
- Parallel I/O
- Dynamic Load Balancing
- QED module
- Particle merging
- OpenMP/MPI/vector parallelism
- CUDA branch/Intel Phi support

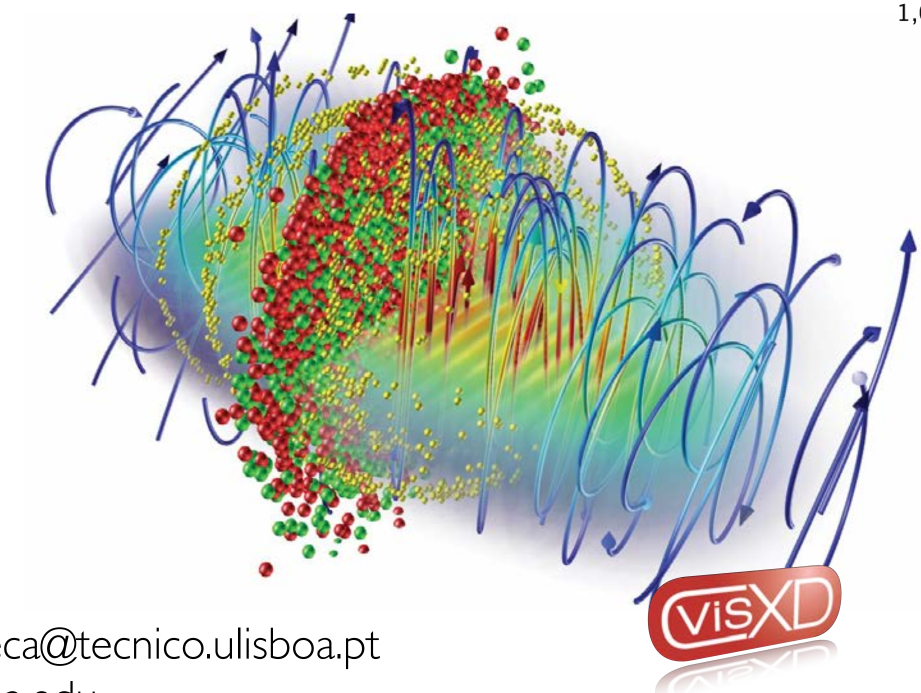

Ricardo Fonseca: ricardo.fonseca@tecnico.ulisboa.pt
Frank Tsung: tsung@physics.ucla.edu

http://epp.tecnico.ulisboa.pt/
http://plasmasim.physics.ucla.edu/

Livingston Curve for Accelerators --- Why plasmas?

Plasma Wake Field Accelerator(PWFA)

A high energy electron bunch

Laser Wake Field Accelerator(LWFA, SMLWFA)

A single short-pulse of photons

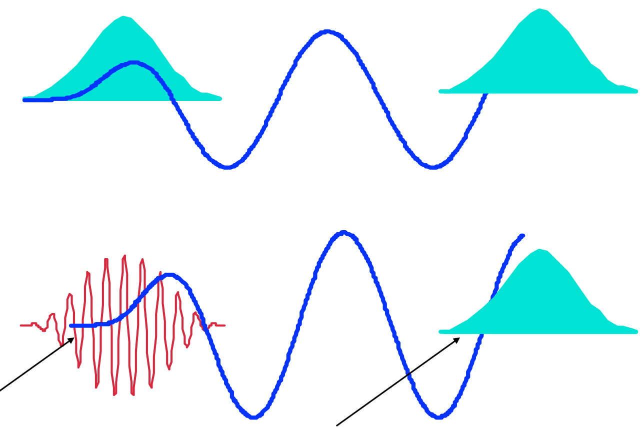

Drive beam

Trailing beam

The Livingston curve traces the history of electron accelerators from Lawrence's cyclotron to present day technology.

Currently plasma based accelerator can match conventional accelerators in terms of energy with much shorter distance. In 2007, the PWFA experiment at SLAC showed energy doubling using 1 meter of plasma.

The goals of our research is no longer to match conventional accelerators in terms of energy, but in terms of quality as well.

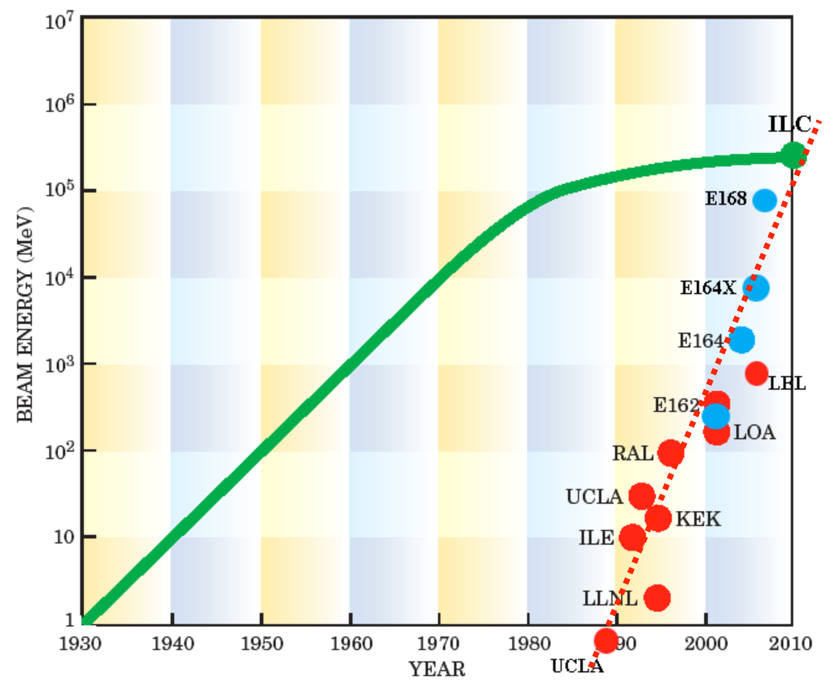

UCLA

X-ray FEL — Coherent light source at Angstrom scale — Can we make compact radiation sources for nuclear science? Using Plasmas?

One application of convention accelerator is a light source. The SLAC accelerator is now a light source called LCLS. In an X-ray FEL (XFEL), a "coherent" electron beam enters an undulator and a bright x-ray comes out, the electron beam can be diverted via an magnet (see right).

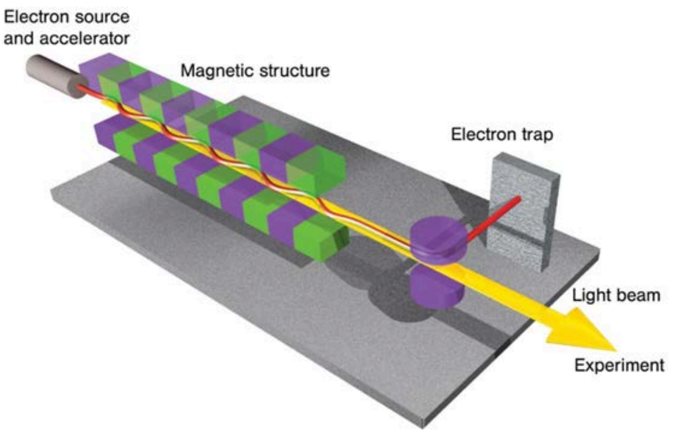

The need for XFEL's light sources can be justified by looking at the light sources in terms of photon energy and "brilliance". Brilliance, also called brightness, is a measure of the coherence of the photon beam (or roughly the # of photons per volume). Improving the brilliance of the beam means the laser light is tightly focused in a small spot, with a very short time duration. This allows the light source to capture very fast phenomenon in a very focused region to study chemical or biological behaviors on a very short (usually femto-second) timescale.

Compared to synchrotron sources, LCLS, which began in 2009, represents a 9 order of magnitude jump in brightness compared to synchrotrons. XFEL's for the first time allow us to probe materials on the nuclear (Angstrom) length scale with femto-second resolution. Laser, while provides high peak brilliance, operates in the ~micron range, which cannot resolve effects on the the nuclear length scale

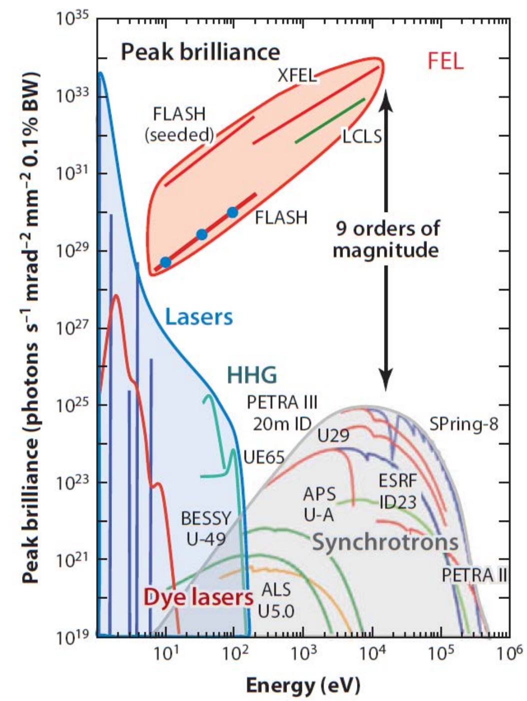

Using PIC simulations, we are trying to study ways to generate high qualities electron beams with high energy and high quality to produce 20keV (0.62 Angstrom wavelength) lights comparable to those generated at LCLS. The beam parameters in LCLS is:

$$\gamma_{beam} = 32,000 = 16GeV$$

peak current density $\sim 0.4 \times 10^{17} \mathrm{A/m^2/rad^2}$

energy spread $\sigma_\gamma/\gamma < \rho \sim 0.001$

What's new this year?

Last year we demonstrated the possibility of using a two electron bunches to double the energy of the witness bunch and produce x-ray comparable to those @ LCLS.

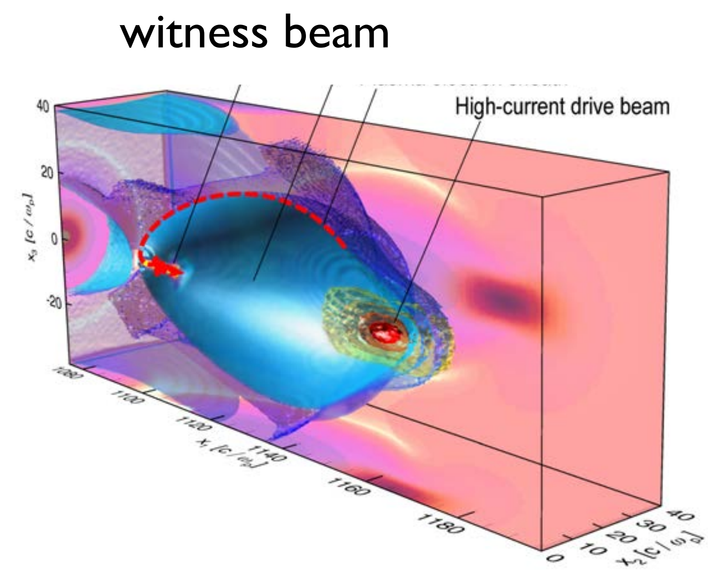

2017

This year we use our numerical tools to study the possibility of generating coherent x-ray using LWFA's in the self-injected regime, where the electrons resonates with the plasma wave near the speed of light. 3D simulations have demonstrated a technique to generate high quality electron beams without an external injector.   (This means that these experiments can be performed without an accelerator) This work was published in late 2017.

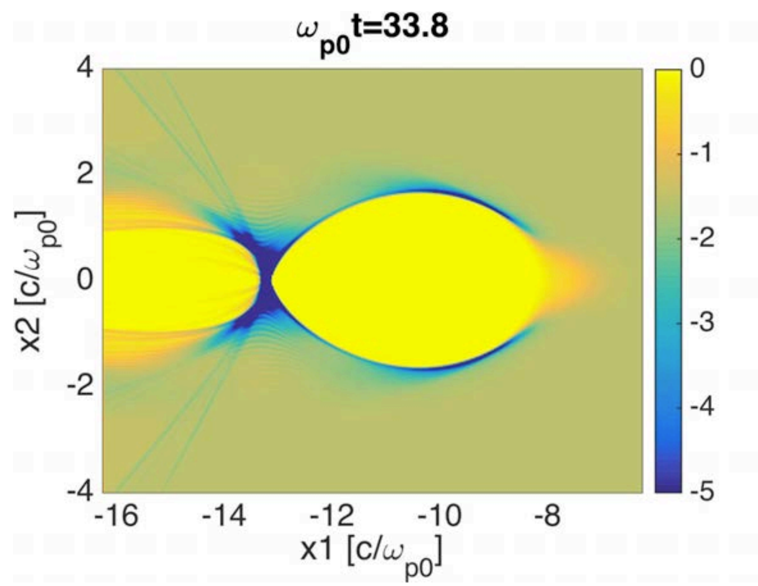

2018

Introduction – Downramp Injection (X. Xu, PRSTAB, 20, 111303 (2017))

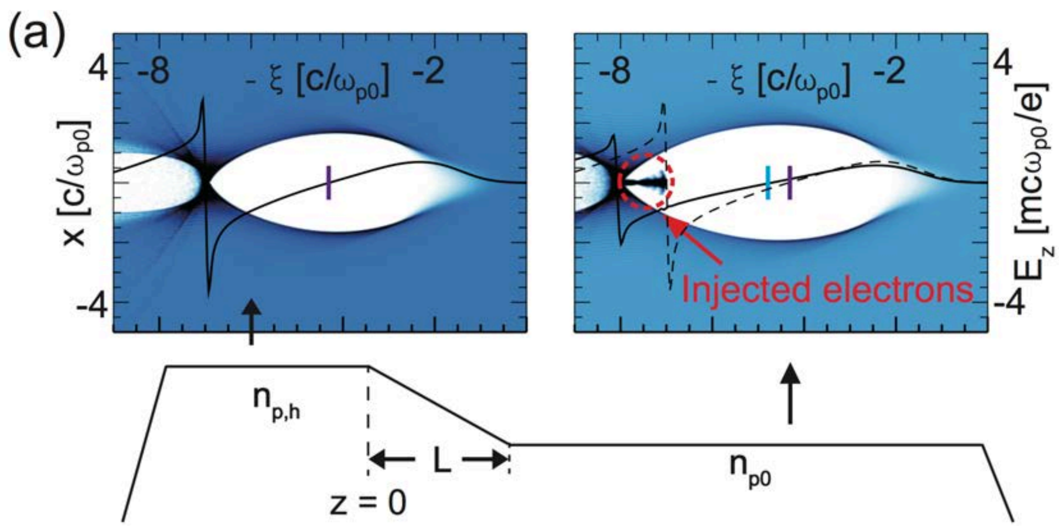

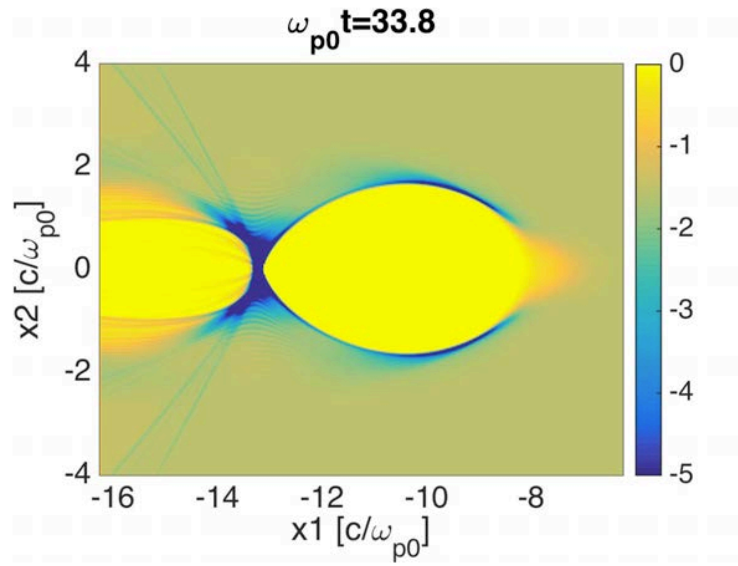

- S. Bulanov[2] et al. (1998), and H. Suk[3] et al. (2001) studied the injection process using **1D** analysis.

[1]T. Katsouleas, Phys. Rev. A 33, 2056 (1986); [2]S. Bulanov, et al., Phys. Rev. E 58, R5257 (1998); [3]H. Suk, et al., Phys. Rev. Lett. 86, 1011 (2001);

Simulation Parameters:

- ~ 1 billion grids in 3D
- 8 particles per cell
- final beam energy from 500MeV to ~GeV, each simulation takes 1 million CPU hour on BW. (3.3mm in this case)
- special EM solvers to eliminate numerical Cerenkov radiation.

	$n_{p,h}$ [cm^{-3}]	n_{p0} [cm^{-3}]	L_{ramp} [mm]	L_{acc} [mm]	Initial T [eV]
Plasma	1.5e18	1e18	1.33 (250 c/ω_{p0})	3.3	10

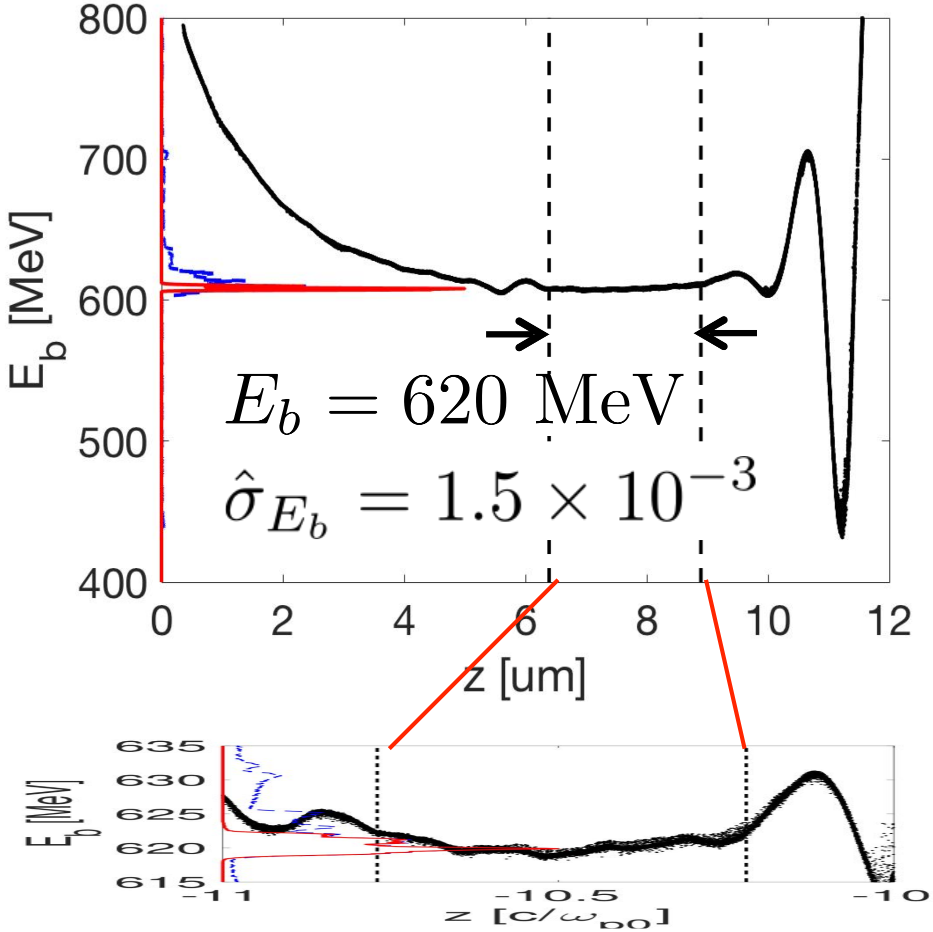

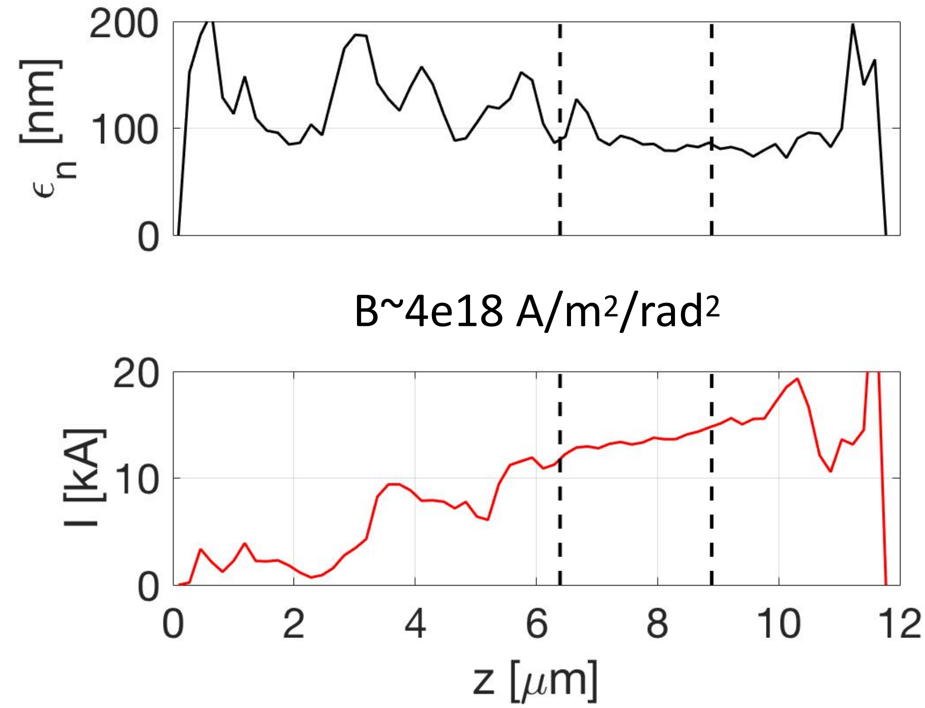

Laser Plasma Interactions in IFE

- IFE (inertial fusion energy) uses lasers to compress fusion pellets to fusion conditions. Inside the fusion chamber (hohlraum), the laser can excite plasma waves and undergo LPI (laser plasma interaction). In this case, the excitation of plasma waves via LPI is detrimental to the experiment in 2 ways.
  - Laser light can be scattered backward toward the source and cannot reach the target
  - LPI produces hot electrons which heats the target, making it harder to compress.

- The LPI problem is very challenging because it spans many orders of magnitude in lengthscale & lengthscale
  - The spatial scale spans from < 1 micron (which is the laser wavelength) to milli-meters (which is the length of the plasma).
  - The temporal scale spans from a femto-second(which is the laser period) to nano-seconds (which is the duration of the fusion pulse). A typical PIC simulation spans ~10ps.

Laser Plasma Interactions

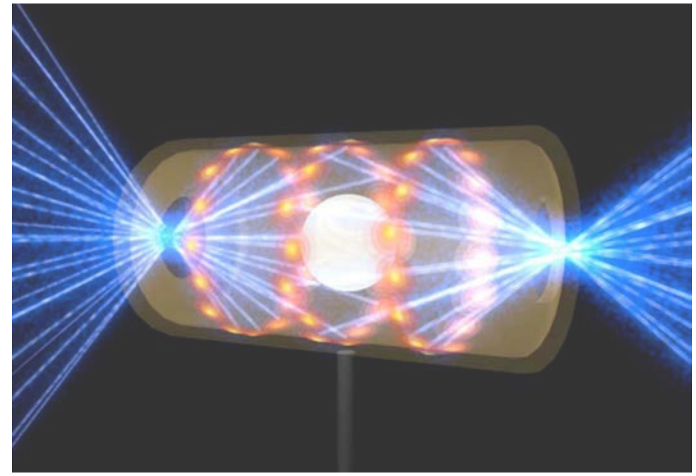

NIF
National Ignition Facility

Lengthscales

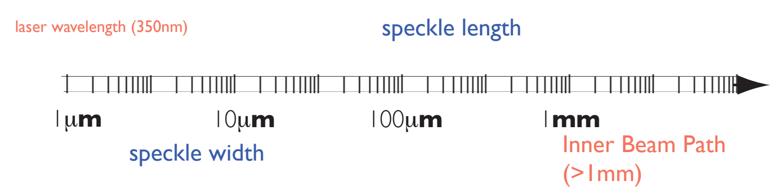

Timescales

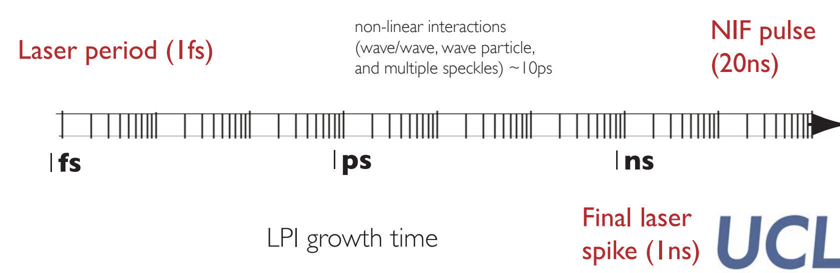

We have simulated stimulated Raman scattering in multi-speckle scenarios (in 2D)

- Although the SRS problem is 1D (i.e., the instability grows along the direction of laser propagation). The SRS problem in IFE is not strictly 1D -- each "beam" (right) is made up of 4 lasers, called a NIF "quad," and each laser is not a plane wave but contains "speckles," each one a few microns in diameter. These hotspots are problematic because you can have situations where according to linear theory, the "averaged" laser is LPI unstable only inside these "hotspots" (and the hotspots can move in time by adding colors near the carrier frequency). And the LPI's in these hotspots can trigger activities elsewhere. The multi-speckle problem are inherently 2D and even 3D.

- We have been using OSIRIS to look at SRS in multi-speckle scenarios. In our simulations we observed the excitation of SRS in under-threshold speckles via:
  - "seeding" from backscatter light from neighboring speckles
  - "seeding" from plasma wave seeds from a neighboring speckle.
  - "inflation" where hot electrons from a neighboring speckle flatten the distribution function and reduce plasma wave damping.

- In the past few years we have added both static and moving speckles into the code OSIRIS. 2D OSIRIS simulations show, that given enough temporal bandwidth, LPI's relevant to IFE (both SRS and HFHI) can be reduced.

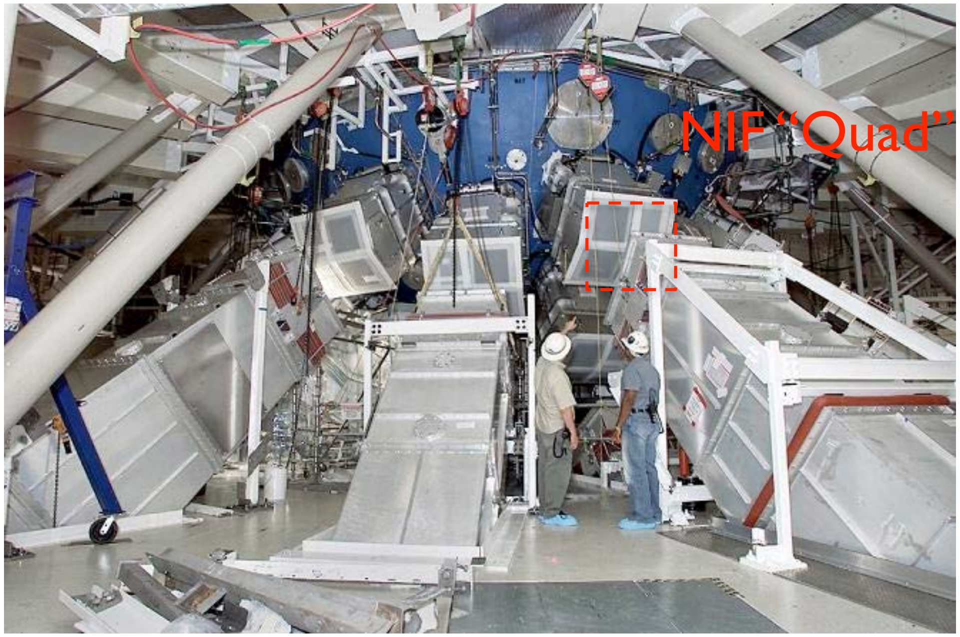

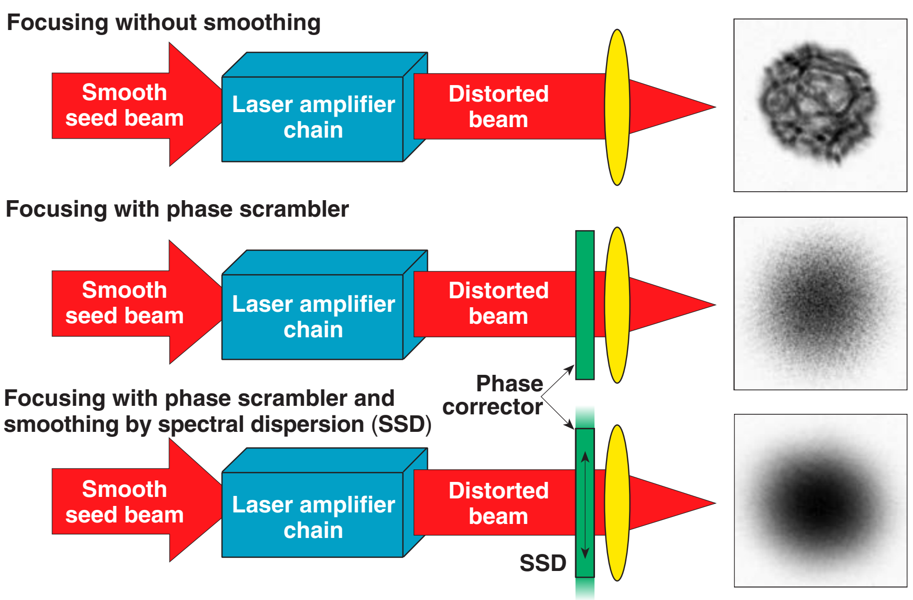

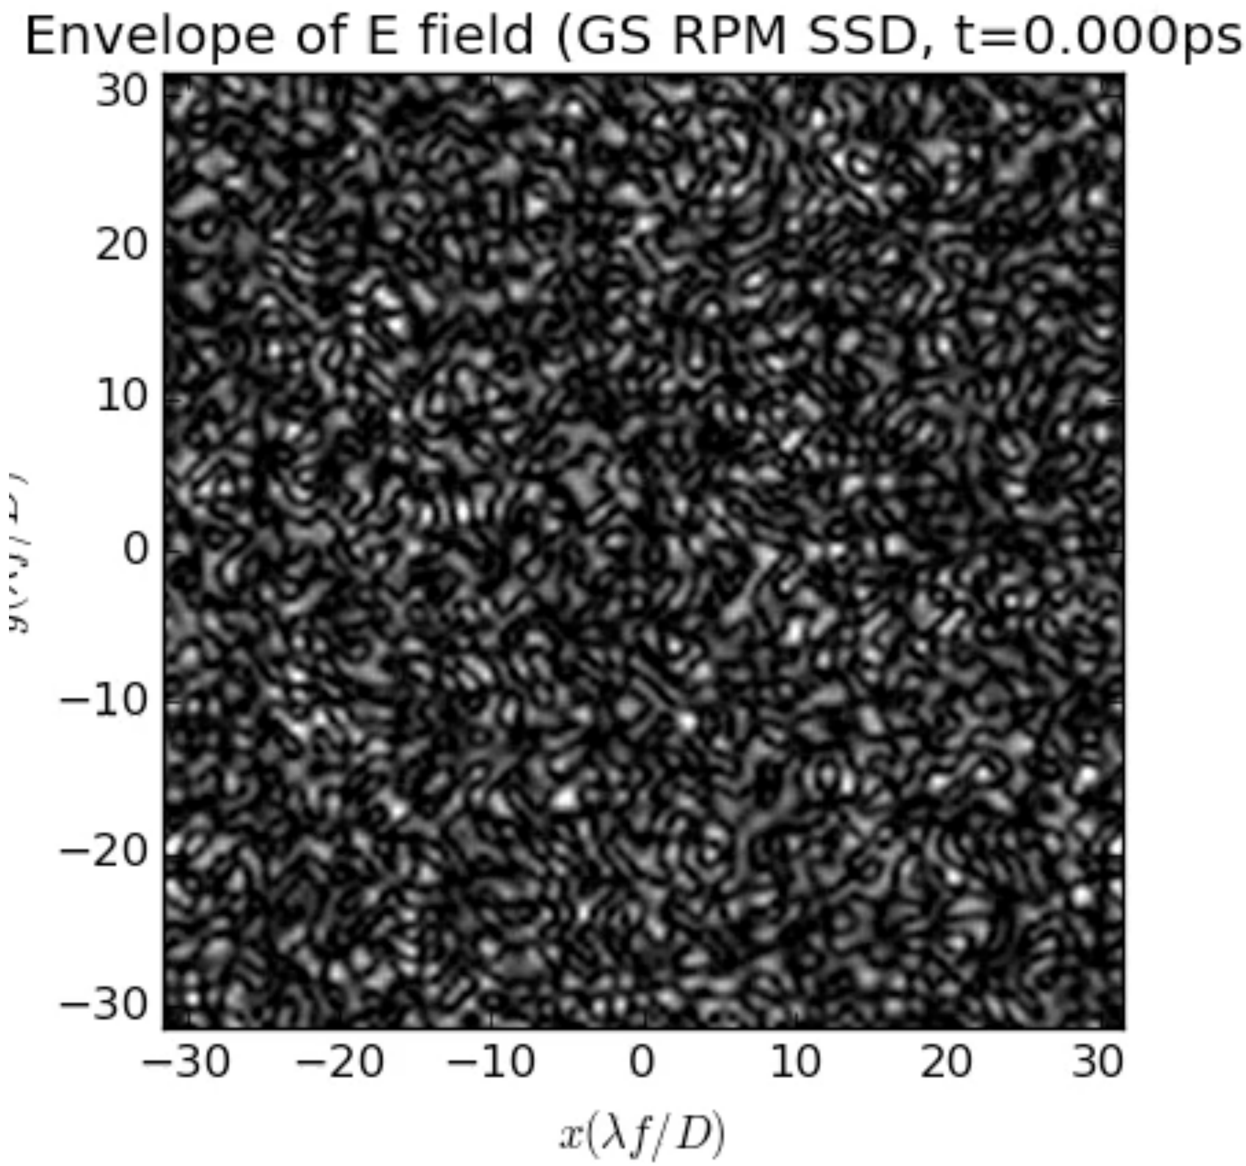

Envelope of E field (GS RPM SSD, t=0.000ps
30
20
10
0
−10
−20
−30
−30
−20
−10
0
10
20
30
$x(\lambda f/D)$

Large scale 2D simulations of SRS with bandwidth (Dr. Han Wen, prepared for publication)

UCLA

Over the past 2 years, we have performed a large number of 2D simulations, ranging from 120 microns to 750 microns long, which is roughly ½ of the total length of the NIF inner beam.

In the past year, we have begun performing simulations with the largest 2D box to date. Typical width of the simulation box is 80 microns, which covers ~28 laser speckles and the typical length is 750 microns (which is > 1/2 of the inner beam path in NIF). Simulations of this scale takes 3-5 million core hours each.

Simulation Parameters:

- T_e= 1-5keV
- density range = 9% to 18% n_c.
- $k\lambda_D$ ~ 0.33 @ z=290 microns.
- laser intensity ~ 1-10 10^{14}W/cm^2

Reflectivities	1D	RPP (f=8)	ISI (1THz)	ISI (6THz)
I_{14} = 5	13%	15%	7%	3%

Linear background density

Immobile ions

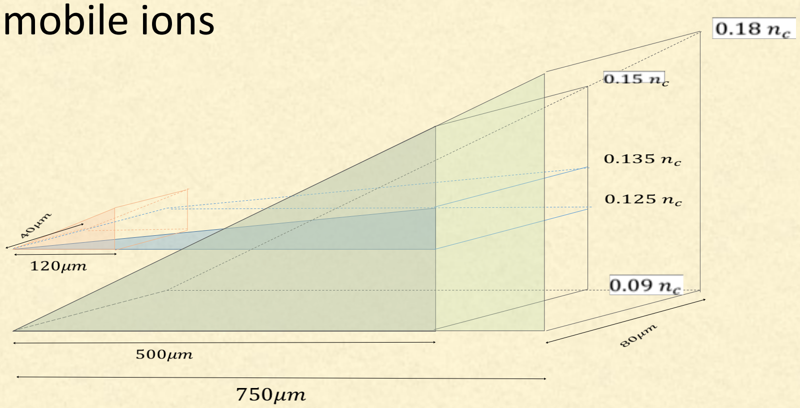

OSIRIS Simulations of multi-speckle LPI with realistic beam smoothing:

RPP **ISI (1THz)** **ISI (6THz)**

longitudinal e-field

transverse e-field

slope $f_e(v)$ near the phase velocity

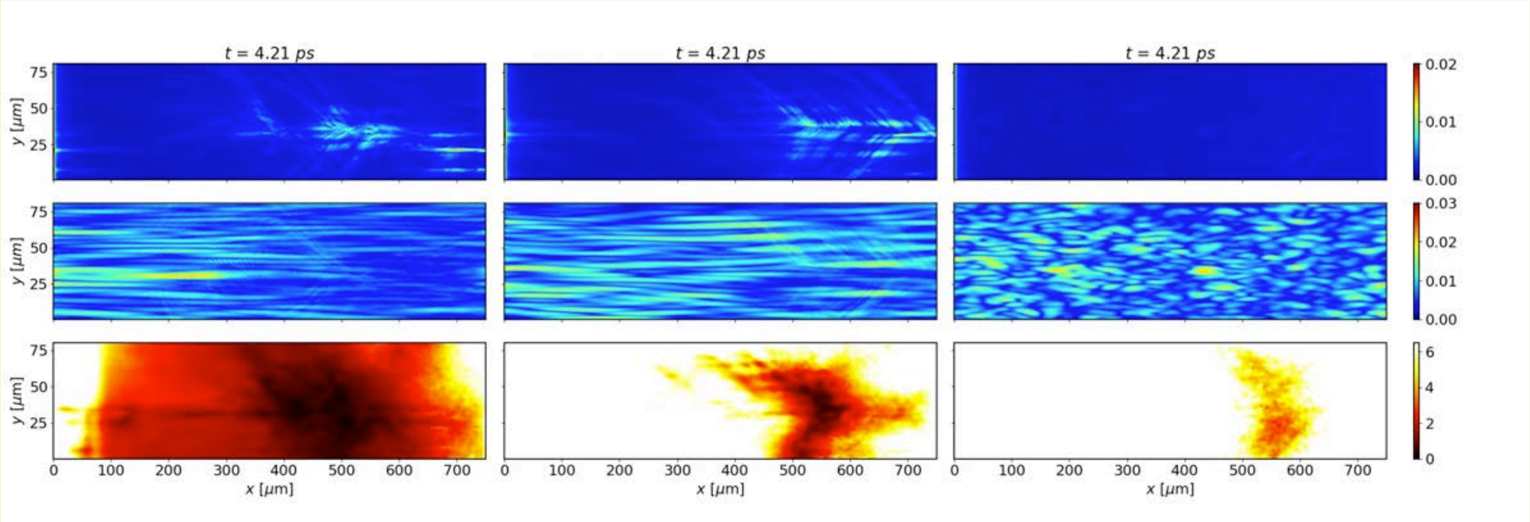

Reflectivities	ID	RPP (f=8)	ISI (1THz)	ISI (6THz)
$I_{14} = 5$	13%	15%	7%	3%

PIC simulations of 3D LPI's is still a challenge, and requires exa-scale supercomputers, this will require code developments in both new numerical methods and new codes for new hardwares

	2D multi-speckle along	3D, 1 speckles	3D, multi-speckle along
Speckle scale	50 x 8	1 x 1 x 1	10 x 10 x 5
Size (microns)	150 x 1500	9 x 9 x 120	28 x 28 x 900
Grids	9,000 x 134,000	500 x 500 x 11,000	1,700 x 1,700 x 80,000
Particles	300 billion	300 billion	22 trillion
Steps	470,000 (15 ps)	540,000 (5 ps)	540,000 (15 ps)
Memory Usage*	7 TB	6 TB	1.6 PB
CPU-Hours	5-10 million	10-15 million	1 billion (2 months on the full Blue Waters supercomputer)

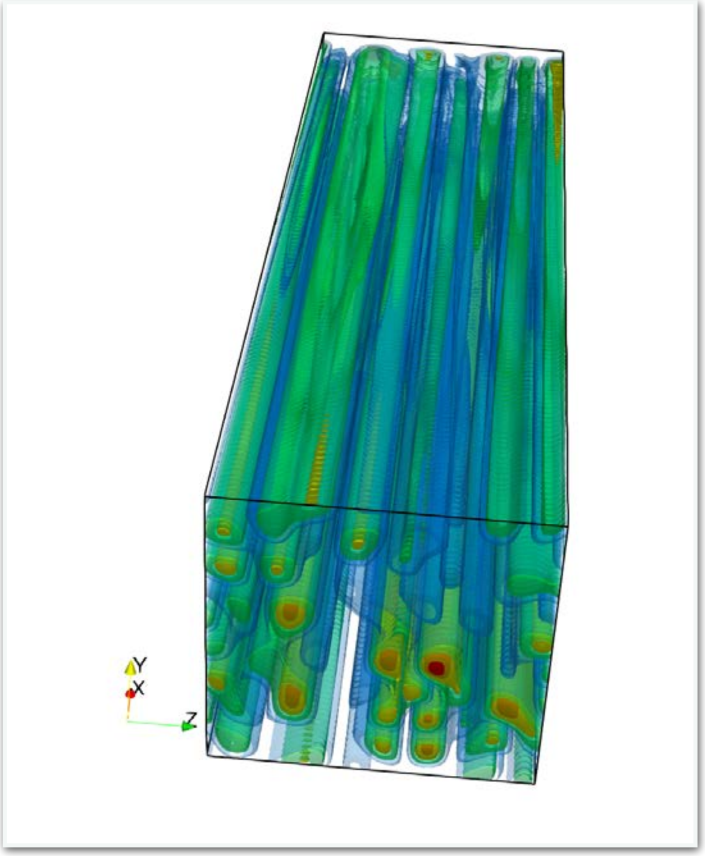

(7 x 7 speckle pattern in 3D produced by OSIRIS)

Designing New Particle-in-Cell (PIC) Algorithms on GPU's

On the GPU (and multi-cores), we apply a local domain decomposition scheme based on the concept of tiles.

Particles ordered by tiles, varying from 2 x 2 to 16 x 16 grid points (typical tile size is 16 x 16 in 2D and 8 x 8 x 8 in 3D)

On Fermi M2090:

• On each GPU, the problem is partitioned into many tiles, and the code associate a **thread block** with each tile and particles located in that tile

We created a new data structure for particles, partitioned among threads blocks (i.e., particles are sorted according to its tile id, and there is a local domain decomposition within the GPU), within the tile the grid and the particle data are aligned and the loops can be easily parallelized.

We created a new data structure for particles, partitioned among threads blocks:

```
dimension part(npmax,idimp,num_blocks)
```

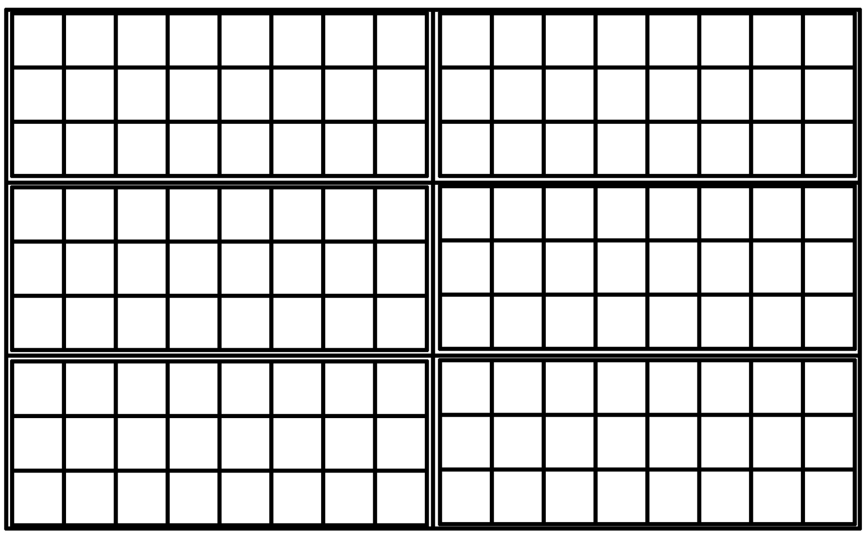

Evaluating New Particle-in-Cell (PIC) Algorithms on GPU: **Electromagnetic Case**

2-1/2D EM Benchmark with 2048x2048 grid, 150,994,944 particles, 36 particles/cell

optimal block size = 128, optimal tile size = 16x16

GPU algorithm also implemented in OpenMP

Hot Plasma results with dt = 0.04, c/vth = 10, relativistic

	CPU:Intel i7	GPU:Fermi M2090	OpenMP(12 CPUs)
Push	66.5 ns.	0.426 ns.	5.645 ns.
Deposit	36.7 ns.	0.918 ns.	3.362 ns.
Reorder	0.4 ns.	0.698 ns.	0.056 ns.
Total Particle	103.6 ns.	2.042 ns.	**9.062 ns (11.4x speedup)**.

The time reported is per particle/time step.
The total particle speedup on the Fermi M2090 was 51x compared to 1 Intel i7 core.

The OpenMP version has been extended to take advantage of of the vector units on the Intel Phi. On the UPIC framework, the particle tasks contains inner loops of length X (where X depends on the particular version of Phi that you are running) and the particles are vectorized automatically by the Intel compiler.

Codes that are described here are available at the **UCLA PICKSC web-site**

http://picksc.idre.ucla.edu/

UCLA

OSIRIS on Intel Phi (Cori supercomputer @ NERSC)

- On the intel Phi, it has multiple levels of parallelization. OSIRIS uses 2 levels of parallelization. Inside each MPI process, particle tasks are vectorized using KNL vector intrinsics. On the Cori supercomputer @ NERSC (1 KNL unit per node, 68 cores per node and 512-bit vector units per core) OSIRIS achieved a speed of nearly 1 billion particles per second on a SINGLE Cori node. (A new version of OSIRIS that incorporates tiling (using OpenMP) is under developed.) Our skeleton code UPIC has 3 levels of parallelism using MPI + Tiles (OpenMP) + automatic vectorization (via Intel compiler) achieved similar numbers, and it is available on the PICKSC website (PICKSC -> Software -> Skeleton Codes -> OpenMP/Vectorization).

- On the Cori supercomputer (which has ~9,500 KNL nodes), OSIRIS achieved nearly ideal weak scaling and excellent (> 90%) strong scaling on nearly the entire machine (8,000 nodes, > 500,000 MPI processes).

- The DOE's first Exa-scale supercomputer, aurora (located @ DOE's ALCF leadership facility), will consist of 500,000 nodes (roughly 50 times the size of Cori). We have applied to be one of 20 teams to use the Aurora supercomputer for 3 months in 2021. This allocation is equivalent to close to one full year of allocation on a current supercomputer and will allow us to model LPI in full 3D.

- Also we are exploring using deep learning as a mechanism to identify regions where kinetic effects are important and use ML to trigger kinetic simulations in 3D on future exa-scale supercomputers.

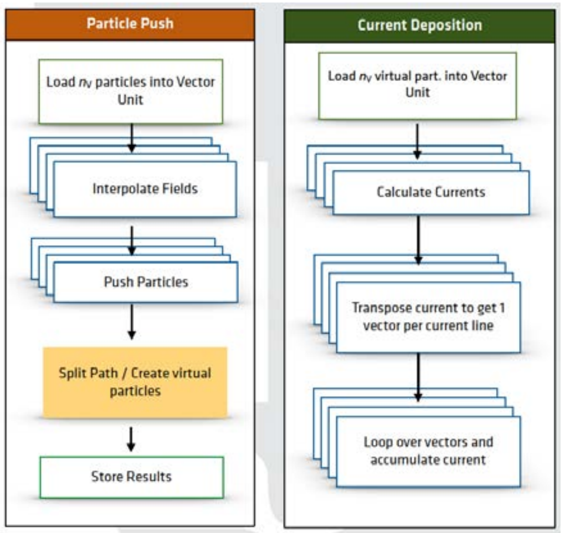

Processing #	Strong Efficiency	Weak Efficiency
1000	100%	100%
2744	95.3%	>99%
4096	95.2%	>99%
8000	90.3%	>99%

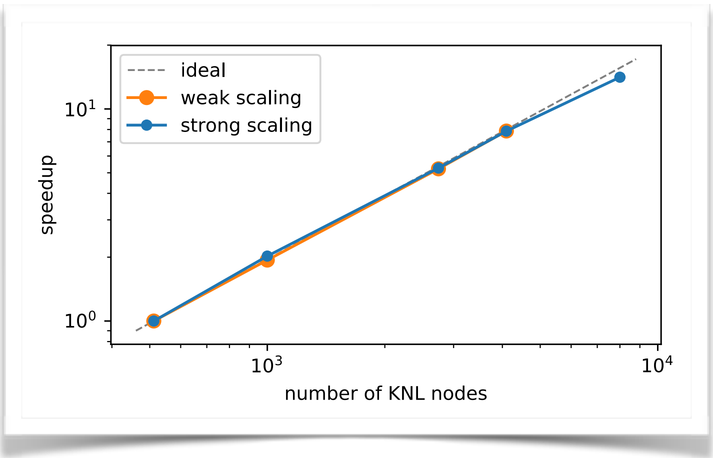

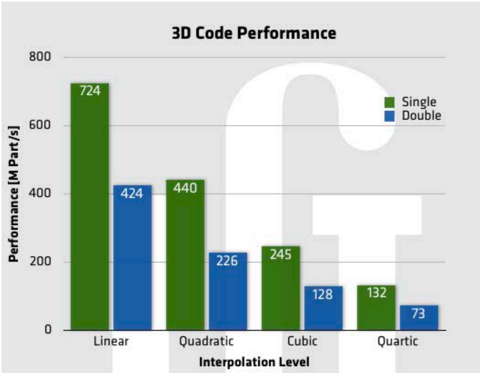

3D Push time [ns]

	Single	Double	Single/Double
Linear	1.38	2.36	58.6 %
Quadratic	2.28	4.42	51.5 %
Cubic	4.08	7.81	52.2 %
Quartic	7.55	13.79	54.7 %

List of publications & presentations in the past 12 months

- Publications:

  X. L. Xu, F. Li, W. An, T. N. Dalichaouch, P. Yu, W. Lu, C. Joshi, and W. B. Mori, "High quality electron bunch generation using a longitudinal density-tailored plasma-based accelerator in the three-dimensional blowout regime", Phys. Rev. Accel. Beams **20**, 111303, 2017

  "Kinetic Simulations of Reducing Stimulated Raman Scattering with Laser Bandwidth in Inertial Confinement Fusion", H. Wen, F. S. Tsung, B. J. Winjum, A. S. Joglekar, W. B. Mori, To be submitted.

  C. Joshi, E. Adli, W. An, C. E. Clayton, S. Corde, S. Gessner, M. J. Hogan, M. Litos, W. Lu, K. A. Marsh, W. B. Mori, N. Vafaei-Najafabadi, B. O'shea, Xinlu Xu, G. White and V. Yakimenko, "Plasma wakefield acceleration experiments at FACET II", Plasma Phys. Control. Fusion 60, 034001 (2018).

  Weiming An, Wei Lu, Chengkun Huang, Mark Hogan, Chan Joshi, Warren Mori, "Ion motion induced emittance growth of matched electron beams in plasma wakefields", Phys. Rev. Lett. 118, 244801 (2017).

- Invited Talks:

  "Petascale kinetic simulations of laser plasma interactions relevant to inertial fusion — controlling laser plasma interactions with laser bandwidth", The 3rd International Conference on Matter and Radiation at Extremes, Qingdao, China, May 2018

- Numerous Presentations, including:

  "Particle-in-Cell Simulations of Laser Plasma Interactions in Multiple Speckles with Temporal Bandwidth", Wen, H., Winjum, B., Tsung, F., et al., 2017, APS Meeting Abstracts, BP11.131

  "Recent progress in simulation and theory toward using nonlinear plasma wakefields to drive a compact X-FEL", Xinlu Xu, Wei Lu, Chan Joshi, W. B. Mori, American Physical Society (APS), Division of Plasma Physics (DPP) conference, Milwaukee, WI, October 23 - 31st, 2017. (Talk)

  Thamine Dalichaouch, Xinlu Xu, Asher Davidson, Peicheng Yu, Weiming An, Chan Joshi, Chaojie Zhang, Warren Mori, "Generating high brightness electron beams using density downramp injection in nonlinear plasma wakefields," American Physical Society (APS), Division of Plasma Physics (DPP) conference, Milwaukee, WI, October 23 - 31st, 2017. (Talk)

  Thamine Dalichaouch, Asher Davidson, Xinlu Xu, Peicheng Yu, Weiming An, Chan Joshi, Chaojie Zhang, Warren B. Mori, "Generating high brightness electron beams using density downramp injection in nonlinear plasma wakefields," Anomalous Absorption (AA), Florence, OR, June 11th - 16th, 2017. (Poster)

Special thanks to Galen Arnold and the Blue Waters staff without whom none of this is possible.